arXiv:1612.03085v1 [physics.soc-ph] 5 Dec 2016

An Epidemiological Mechanistic Model of Box Office Dynamics

Naghmeh Momeni[1,*], Amir Tohidi Kalorazi[2], Michael Rabbat[1], and Babak Fotouhi[3,4]

[1]Department of Electrical and Computer Engineering, McGill University, Montreal, Quebec, Canada
[2]Department of Electrical Engineering, Sharif University of Technology, Tehran, Iran
[3]Program for Evolutionary Dynamics, Harvard University, Cambridge, MA, USA
[4]Institute for Quantitative Social Sciences, Harvard University, Cambridge, MA, USA
[*]naghmeh.momenitaramsari@mail.mcgilll.ca

ABSTRACT

In this paper we propose a mechanistic model that links micro social interactions to macro observables in the case of diffusion of film-going decisions. We devise a generalized epidemic model to capture the temporal evolution of box office revenues. The model adds social influence and memory effects to conventional epidemic models. Fitting the model to a temporal data set containing domestic weekly revenue of the 5000 top US-grossing films of all time, we find that a two-parameter model can capture a remarkable portion of the observed variance in the data. Using the distribution of the estimated parameters for different genres, we then present a predictive model which provides reasonable a-priori estimates of future sale as a function of time.

Introduction

Understanding human collective behavior is a much-studied topic across a diverse array of disciplines. There are retrospective approaches and predictive approaches. In the former, the aim is to understand how (and why) certain collective social movements emerged. Examples include revolutions[1,2], sudden emergence of radical nationwide attitudes[3,4], financial crises[5,6], and even more mundane subjects such as fashion trends and fads[7,8]. The level of complexity of collective social problems is too high to allow simple answers, which is manifest in the fact that for many of the said problems, there is still no convergent framework, and consensus is rarely reached on how (and especially why) certain social phenomena emerged. The predictive approach is more common in disciplines such as public policy, finance, economics, and marketing. The aim is to forecast how people would react to certain policies, technologies, products, etc., with reasonable degree of accuracy. Consumer behavior, word-of-mouth effects and herd behaviors are central topics in behavioral finance and marketing research[9–12]. The main reason is clear; in a socio-economic system whose centerpiece is the free market, investors naturally like to increase investment efficiency by foreseeing the future, even partially. Though progress is being made, the status-quo in many domains is still far from near-certain prediction. The popular adage "Half the money I spend on advertising is wasted; the trouble is I don't know which half", often attributed to the marketing pioneer John Wanamaker[13], is still relevant to the present-day state of knowledge. The popular accounts on this challenge of prediction is also ample. One example is the rejection of *Harry Potter* by 12 publishers before finally getting published, which summarizes the current ability of even experts in predicting product success and consumer behavior.

In this paper, we focus on a particular example of consumption, and seek to describe it with as simple a model as possible. We focus on films. Forecasting box office performance of films is a well-researched topic[15–21]. In the post-Internet era (which has transformed the typical size and scope of

available data), more and more attention is being paid to search counts, weblog content, and social media trends, to understand and predict consumer behavior in every marketing context, including film revenues. Most of these studies aim to predict future revenues based on pre/upon-release data, such as pre-release critiques[22,23], Twitter mentions[24], YouTube views[25], pre-release news mentions[26], Wikipedia page hits and edits[27], and search engine query data[28]. Although studio research teams are advancing knowledge on what makes a film popular, the current understanding of consumer behavior in this context and the current ability to predict future box office success is still far from certainty.

In contrast to the approaches of the studies mentioned above, we take a mechanistic, bottom-up approach in this paper to study the evolution of film popularity. We are interested in modeling the micro interpersonal mechanisms that drive this dynamics, and seek to capture the emergence of macro observable outcomes (which has the advantage of data availability). Films constitute a special kind of commodity, with idiosyncratic features in terms of diffusion. In the present stage of capitalist economy, most products are subject to a constant influx: frequently, new versions of products (e.g., cell phones) arrive, and people update their possessions at varying speeds (some must have the most recent iPhone as soon as possible, some wait longer, some are pushed via planned obsolescence, etc.). This is more or less true for a considerable share of products with use: cell phones, TVs, video game consoles, laptops, and tablets, to name a few examples. Films, however, have different dynamics (let us emphasize that we focus only on box-office sales; we are not considering subsequent revenues via DVD release and other royalties). Almost all people go to the cinema *together*. That is, film-going is a joint activity. So the social component that drives the film-going decision is strong. There is also a time factor: films do not stay on the screen indefinitely and are sequentially replaced by new ones. In this paper we seek to construct a minimal model that captures these features.

We devise a basic model of social influence derived from conventional epidemic models (see[29] for a thorough review of these models) with a twist. By adding memory effects to a basic epidemic-type model, we obtain a two-parameter model which is analytically tractable. We use data on the weekly domestic (in the USA) sales of 5000 recent films to fit model parameters, and show that the model can reliably capture the temporal evolution of the box-office revenue. Although our main aim is to shed light on the micro social mechanisms, and not prediction, we show that as a byproduct of the analysis, the results can also have considerable prediction-oriented utility with remarkable accuracy, which we do analyze but not as the central theme of the paper.

Results

model

We denote an individual who has seen the film with I (for *infected*) and one who has not seen the film by S (for *susceptible*). Social influence results from a person observing friends, family, peers, etc., having seen a film. If the proportion of social ties in one's proximity having seen the film increases, the person is more likely to catch up with the trend. For a given social tie between an S individual and an I individual, we denote the 'transmission' rate by β. That is, for a small time increment dt, the probability that the S individual gets infected as a result of this social tie is βdt. It is worth noting that in actuality, there are two distinct compartments that both exert influence on an S individual to see the film. One consists of social ties who have already seen the film, and the other consists of those who have decided to see the film but have not done so yet. Since people rarely go to the theater alone, the latter group is expected to have considerable contribution in the social influence on an individual. Because model simplicity is our first priority in this study, we approximate the situation by assuming that the I compartment consists of both of the subpopulations mentioned above, and we assume that β is an average transmission coefficient for

both types. In other words, we assume that 'having decided to see the film' and 'having seen the film' are interchangeable, and we are ignoring the cases where a decision is made but eventually not followed through on.

Another central element of the dynamics is memory. There is a constant inflow of new films being released. Human cognitive capacity is limited. The 'buzz' created by a newly-released film inevitably decreases as time advances. A simple choice to model human memory is an exponential function e^{-Bt}, where $B > 0$ is a constant[30,31]. The larger the value of B is, the shorter the lifetime of the film will be in the collective memory of the society. Also, we expect film lifetime (the number of weeks theaters screen it) to decrease with B. The reason is that increasing B makes the film being forgotten more quickly, and theaters only keep screening a films until the generated revenue matches the costs.

Let us denote the cumulative box office revenue at time t by $G(t)$, and let ρ_t denote the proportion of individuals in the I compartment at time t. The temporal evolution of $\rho(t)$ under the mean-field regime is given by

$$\dot{\rho}_t = (1-\rho_t)\rho_t \beta \langle k \rangle e^{-Bt}, \tag{1}$$

where $\langle k \rangle$ is the average number of social contacts of individuals. Denoting $\beta\langle k \rangle / B$ by A, and noting that $G(t)$ is proportional to ρ_t, we obtain the following expression for $G(t)$:

$$G(t) = G(0)e^{A(1-e^{-Bt})}. \tag{2}$$

It is worth noting that A models social influence, that is, stronger A means stronger effect from the social surroundings on the decision of filmgoers. The value of A can increase in two ways: increase in the connectivity of people (through $\langle k \rangle$), or increase in the per-link transmission rate.

An alternative approach to modeling social influence is taking into account group pressure, instead of individual influence. For example, in the conventional voter model (which is the basic model of opinion dynamics with a vast literature of studies and extensions[34]), the probability of switching to one state from another is determined by the respective proportions of the states at the individual's vicinity. In the epidemic approach, we assumed that each social tie exerts an influence to an individual which is independent of the number of social ties that person has. An S person with three I friends receives the same amount of influence whether this person has 20 friends or only has those three friends. The voter model takes a different approach, and considers proportions instead of absolute numbers. This means that the S person with three I friends will receive an influence proportional to $3/20$ if the person has 20 friends, and proportional to $3/3$ if the person has three friends. An alternative voter-type model can be devised for the problem at hand which leads to the same end result under the mean field approximation. If individual x with state S has k_x friends and I_x of them have state I, then the probability that individual x will change its state to I is given by $\alpha dt(I_x/k_x)e^{-Bt}$. The memory factor is identical to the previous scenario, because memory-wise there is no difference between the two cases. The factor α is just the rate of strategy revision, that is, how frequently on average do individuals consider updating their state. In the methods section we show that in this case, The temporal evolution of $\rho(t)$ under the mean-field regime is given by

$$\dot{\rho}_t = (1-\rho_t)\rho_t \alpha e^{-Bt}, \tag{3}$$

which is identical to (1), with only the name of the coefficient changed. The parametric form of equation for $G(t)$ will be identical to the one obtained above.

Parameter Estimation

Let us define $z(t) \stackrel{\text{def}}{=} \log\left[G(t)/G(0)\right]$. To estimate A and B in (2) for a given film, what we essentially need to do is to fit the following model to empirical time-series data of film sales: $z(t) = A(1-e^{-Bt})$. Minimizing the sum of squared errors (denoted by E) via conventional algorithms such as gradient descent or Newton's method is then straightforward, especially because the gradient and the Hessian are very simple to calculate:

$$\nabla E = \sum_t \begin{bmatrix} -2\left(1-e^{-Bt}\right)\left(z_t - A\left(1-e^{-Bt}\right)\right) \\ -2Ate^{-Bt}\left(z_t - A\left(1-e^{-Bt}\right)\right) \end{bmatrix}, \quad H = \sum_t \begin{bmatrix} 2\left(-1+e^{-Bt}\right)^2 & e^{-2Bt}\left(2te^{Bt}(2A-z_t)-4At\right) \\ e^{-2Bt}\left(2te^{Bt}(2A-z_t)-4At\right) & 2At^2\left(2A+e^{Bt}(z_t-A)\right)e^{-2Bt} \end{bmatrix} \tag{4}$$

But to make the proposed procedure more practicable for a broader readership and further emphasize the strength of the model, we take an even simpler approach. With an approximation, we transform the present problem into one of simple linear regression. We achieve considerable simplicity at the cost of losing some accuracy, which is reasonably small (as will be demonstrated below). Note that in the limit as $t \to \infty$, we have $z(t) \to A$. This means that if a film would be allowed to run forever, $z(t)$ would approach A. As an approximation to infinite time, we take the up-to-last week revenue of each film as an estimate of A. Denoting the lifetime of film by L, this means that we can estimate A by $z(L)$. So we need to estimate B in the equation $z(t) = z(L)(1-e^{-Bt})$, which essentially means that with the transformation $y(t) \stackrel{\text{def}}{=} \log\left[1 - z(t)/z(L)\right]$, we have an simple linear regression problem in which $y(t)$ is a linear function of t, and B can be readily estimated for each film. The histogram of the regression R^2 for the 5000 films in the data set is presented in Figure 1 as stratified by (a) genre, (b) month of release, and (c) year of release. It can be readily seen that the increase in error in the said simplification has not been substantial. The histogram of estimates for A is presented in Figure 2 as stratified by (a) genre, (b) release month, (c) release year. The histogram for B is presented in Figure 3 with similar stratification.

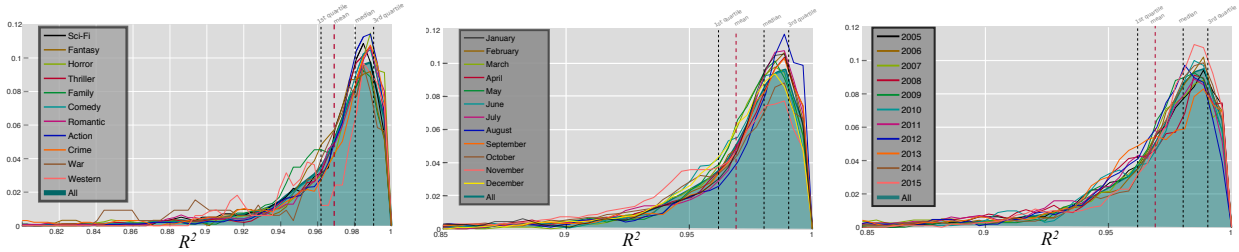

Figure 1. The histogram of adjusted R^2 for films as categorized by (a) genre, (b) month of release, (c) year of release.

Figure 4 depicts A and B as a function of the number of weeks the film remained in theaters (L). It can be seen that A has a positive relationship with L, and that B has a negative relationship, as expected (because increasing A—social influence—means increasing revenue, and increasing B means faster memory decay).

Denoting the estimated revenue from the above procedure by $\widehat{G}(t)$, we can plot $\widehat{G}(t)/G(t)$ to see how accurate the prediction is. Because of different lifetimes of different films (that is, different L values due to different number of weeks films survive in theaters), we plot it as a function of normalized time, which is defined as t/L, so that we can plot all films at the same frame. Figure 5a presents the results, stratified by

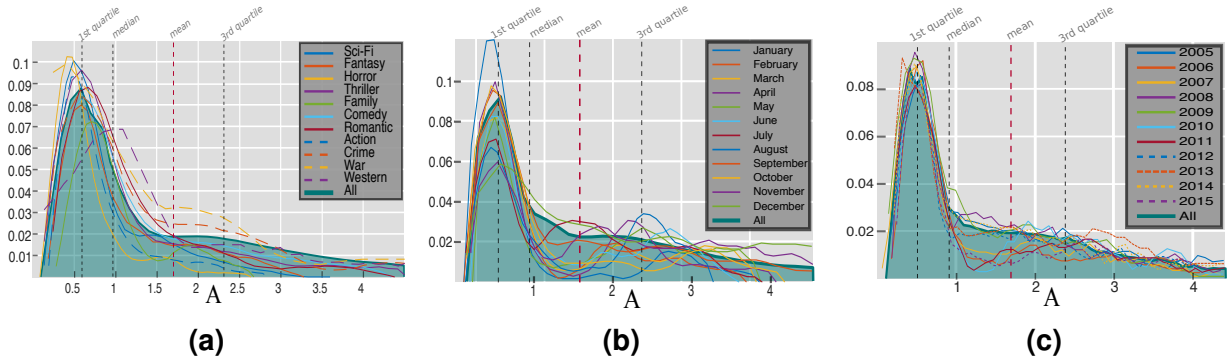

Figure 2. The histogram of A (parameter of model as given by equation (2)) for films as categorized by (a) genre, (b) month of release, (c) year of release.

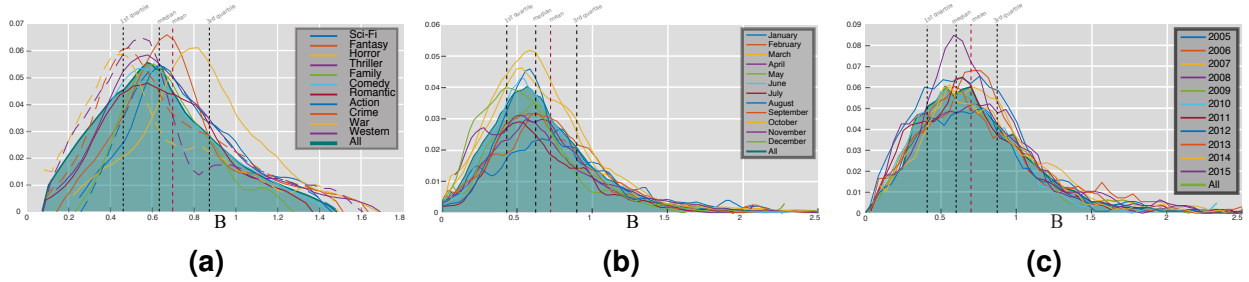

Figure 3. The histogram of B (parameter of model as given by equation (2)) for films as categorized by (a) genre, (b) month of release, (c) year of release.

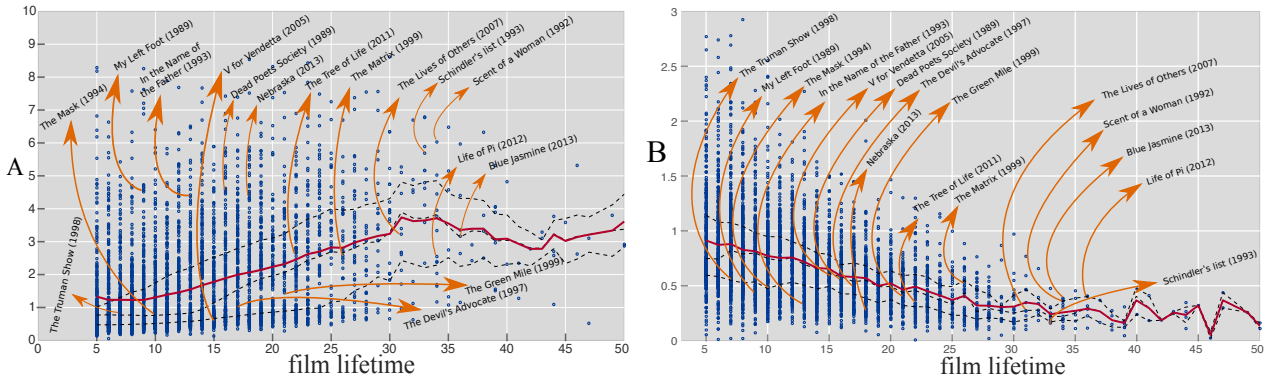

Figure 4. (a) and (b): Model parameters as a function of film lifetime, and (c): A as a function of $\frac{1}{B}$. The red line pertains to the mean, and the dashed lines pertain to the 25, 50, and 75 percentile. Some example films are pointed out for illustrative purposes.

genre. Given the model simplicity, the errors are reasonably low. Moreover, there is no visually discernible difference between the performance in different genres.

We also attempted to simplify the model even further by eliminating one of the parameters by well-approximating it as a function of the other parameter, so that we would have an elegant one-parameter model. Visually, we find that for a considerable portion of the films, the relationship of A and $1/B$ seem close to linear, as depicted in figure 5b. The correlation of A and $1/B$ is indeed high ($\rho = 0.60$). But

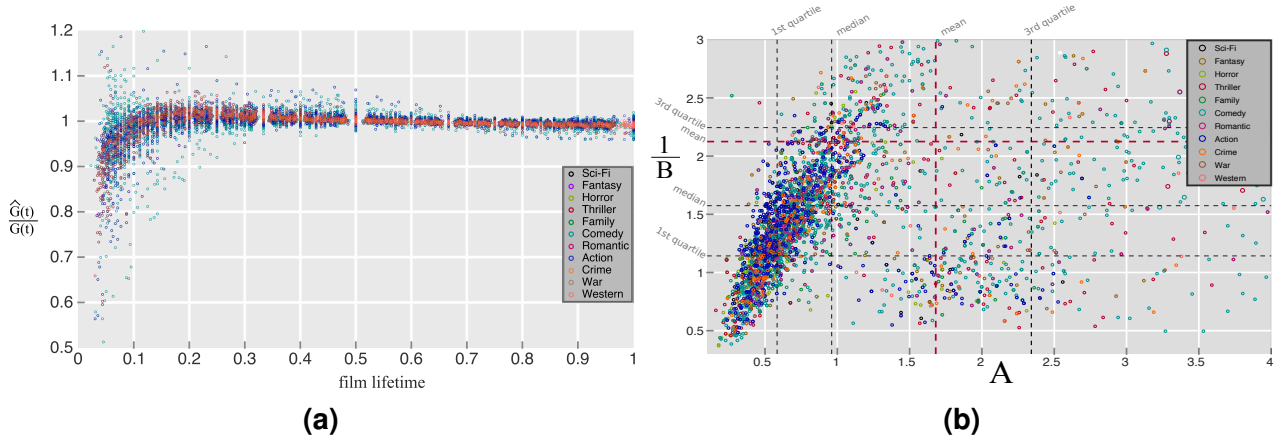

Figure 5. (a) the ratio of the predicted revenue over actual revenue for all films as a function of normalized time, that is, t/L. (b) A as a function of $\frac{1}{B}$

since the linear relationship seems to hold only for a certain portion of the population, which is indeed the majority but not a strong one, we refrain from complete replacement of one of the parameters in terms of the other one based on this relationship. We also attempt such a simplification by investigating possible linkages between A and B to G_0 (the first-week film revenue). Figure 6 presents the scatter plots of A and B as a function of G_0. There is a clear negative relationship between A and $\log G_0$ ($\rho = -0.59$, and $p < 10^{-5}$ for the t-test of the linear-association hypothesis). There is also a small positive association between B and $\log G_0$ ($\rho = 0.22$, and $p < 10^{-5}$ for the t-test of the linear-association hypothesis). The association is weak because B is modeling human memory in our simple model, so it is not expected to vary much with other parameters.

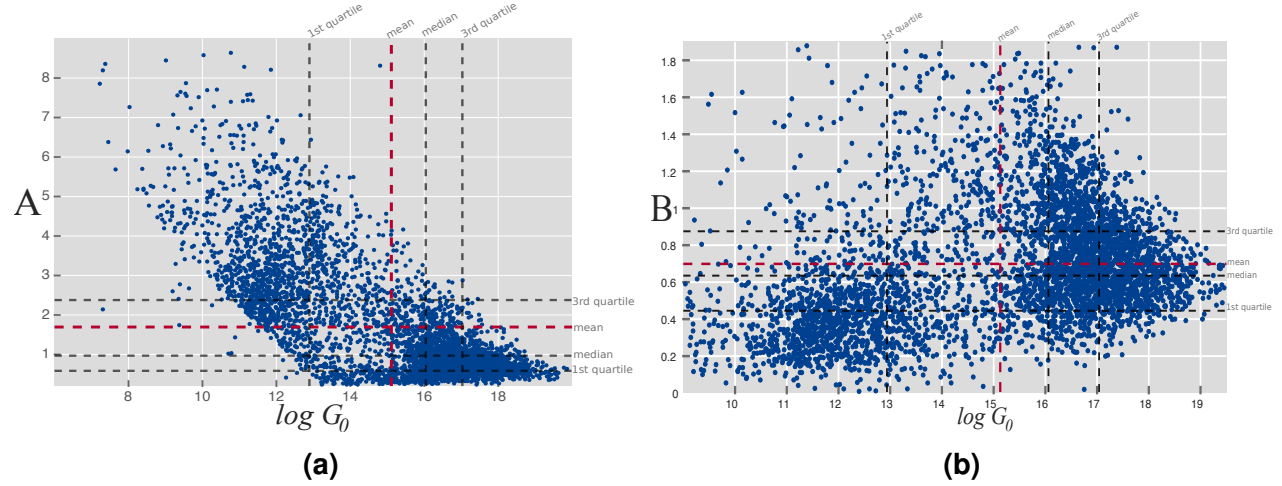

Figure 6. The parameters of the model as a function of the revenue in opening week. The red line pertains to the mean, and the dashed lines pertain to the 25, 50, and 75 percentile.

As we emphasized above, the main focus of this study is to devise a mechanistic model that links the micro mechanisms of social influence to macro observable phenomena, and sales prediction is not a primary goal. But as a byproduct, we can utilize the results in order to add a predictive dimension to the results. We have estimated a value for the A parameter and one for the B parameter for each film. To obtain a formula that would function as a reasonable 'average behavior' for every film, we can

take two approaches. The first approach is to pool all the A estimates together and all the B estimates together, and take the medians as the parameters for the predictive model. The other is to first stratify by genre (or any other attribute, we take genre here as an illustrative example), and take the median of the distributions of parameters only for films in that genre to predict sales of films in each genre. Figure 7 presents the results in these two cases. It is seen that the pooling procedure produces better results than the genre-specific method. Also note that conventional cross-validation methods are not applicable here, because in estimating the parameters for each film, other films have not been used. In other words, our data units are time series here, not data points one would fit a curve to. So cross validation would simply act on the median-taking procedure, not the parameter-learning process.

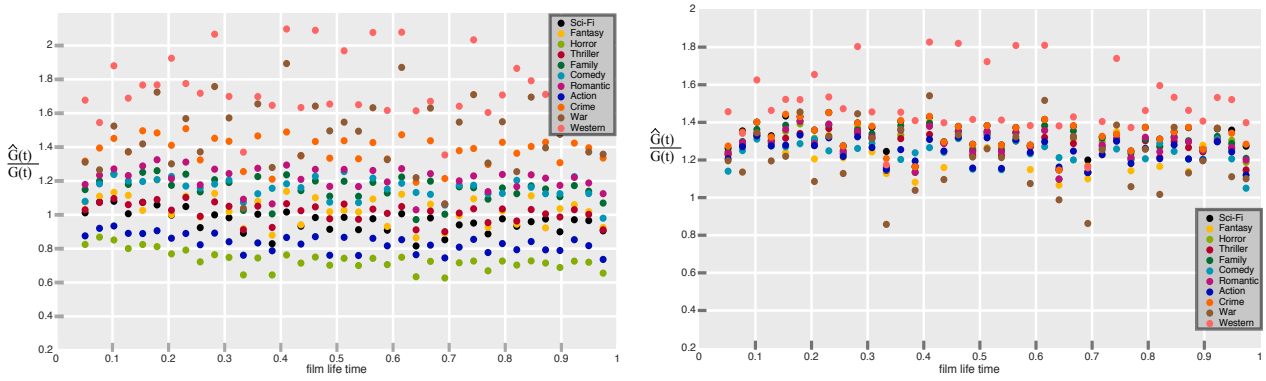

Figure 7. Ratio of predicted value of $G(t)$ to the empirical value as a function of normalized film lifetime, using median values in the distributions of A and B across (a) all films pooled together, (b) only films of given genre.

We also investigate if G_0 can be replaced by pre-release data. As an illustrative example, we use the Wikipedia page hits (W) of films in the month prior to release as a proxy for G_0. As Figure 8 presents, there is a positive association between $\log G_0$ and $\log W$ (with $\rho = 0.44$). Accurate prediction of $G(0)$ from pre-release data has been already done in the literature with remarkable accuracy (e.g., $R^2 > 0.9$ using Wikipedia page edits and hits[27], and $R^2 > 0.95$ using Twitter mentions one day prior to release[24]). So for the reader interested primarily in prediction, there is an avenue for improvement of prediction ability by combining pre-release data with the results presented in this paper.

Discussion

The distribution of A (presented in Figure 2) is skewed. It is known in the literature that film revenues have a heavy-tailed distribution[32,33]. Since greater A means greater social influence, and this in turn means more revenue (and longer lifetime in the theaters, as shown in Figure 4a), we would expect that such a heavy-tailed distribution would be reflected in the distribution of A, too. As we discussed above, B (that characterizes memory decay) models intrinsic properties of the human memory system and is not expected to vary much across films. As Figure 3 illustrates, B is fairly localized, which does match this expectation. Unlike A, parameter B is expected to have a negative relationship with revenue, and consequently, with L (film lifetime). This is confirmed empirically, as illustrated in Figure 4b). The expected increase of L with A and its decrease with B are both statistically significant (hypothesizing a linear relationship, we have $p < 10^{-5}$ for both t-tests).

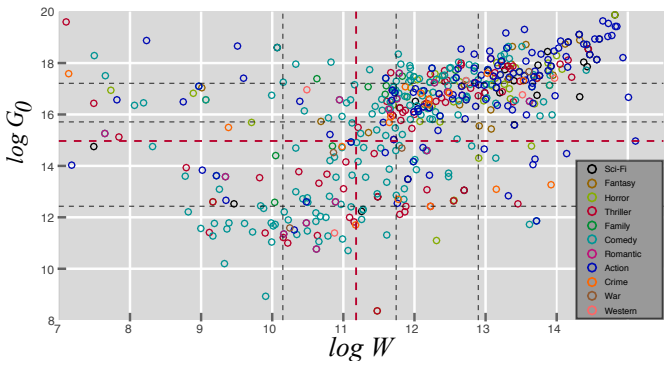

(a)

Figure 8. The parameters of the model as a function of the revenue in opening week

As illustrated in Figure 6, there is a clear negative relationship between A and $\log G_0$ ($\rho = -0.59$, and $p < 10^{-5}$ for the t-test of the linear-association hypothesis). This is expected because a higher G_0 means a stronger role for marketing campaigns and the direct influence individuals receive by encountering advertisement. For example, multitudes of people were individually waiting for the release of the recent Star Wars: The Force Awakens (2016) movie, not waiting for the feedback of others to make a decision.

There is also a small positive association between B and $\log G_0$ ($\rho = 0.22$, and $p < 10^{-5}$ for the t-test of the linear-association hypothesis). The causes of this positive association are purely structural. It does not mean that movies with stronger premier are forgotten faster. It merely reflects the fact that the cinema market capacity is limited. That is, people are either 'movie-goers' or not, and sudden change of attitudes towards the film industry is not very common, especially in a single film lifetime. So higher G_0 necessarily means that the film has to decay faster. For example, Finding Dory (2016) has opening week revenue 231M, and The Jungle Book (2016) has 130M, but both of them sold 31M in the 4th week. It is in this sense that we say the one with higher G_0 has to decay faster.

The predictions of the aggregated model are more accurate when all films are pooled (Figure 7a) than when genre-specific pooling is used (Figure 7b). The accuracy differs for different genres. It is of note that genres with similar content do exhibit similar performances. War is close to Western, Romantic is close to Family as well as Comedy (primarily due to Romantic-Comedy films that bring these two genres close), and Sci-Fi is close to Fantasy. Also, note that differences in accuracy are not due to different subpopulation sizes (that is, it is not the case that, for example, performance for Sci-Fi is good because there are many Sci-Fi films in the data set, and that of War is poor because there are only few War films), because when genre-specific pooling is used (Figure 7b), the performances are still distinct. This means that, as one would intuitively expect, there exist inherent differences between films of different genres, particularly with regard to our model.

Since the model presented in this paper is minimal and we insisted on model simplicity as the central motive of this study, extensions to the present model can readily be made. Which direction yields the most improvement is an interesting question on its own. We ignored individual influence received from marketing campaigns, which can be added to the model with a new parameter. That way, there will be two distinct sources of influence. One would be social, as included here, and the other will be individual (through direct contact with advertisements in the media, billboards, etc.). The latter will add a new parameter to the model. One conjecture would be that susceptibility to social influence would be highly correlated with susceptibility to direct marketing, so that the new parameter could be reasonably well-approximated as a function of already-existing parameters. Another improvement would be to use data

from social media to incorporate the effects of network position of individuals on their social influence, and the effects of structural properties of the underlying social networks in the diffusion of decisions.

Methods

Model Analysis

Let us denote the state of individual x by s_x, which is zero if the individual is in the S state and is one if the individual is in the I state. For a given I social tie, the probability that after time dt the state will be transmitted to individual x is given by βdt, by definition. The probability that transmission does not occur is $1-\beta dt$. The probability that the individual does not get infected from any of existing social contacts is $(1-\beta dt)^{i_x}$, where i_x is the number of social contacts of individual x who are in state I. Thus, the probability that individual x will be infected after time dt is $1-(1-\beta dt)^{i_x}$, which can be Taylor-expanded to first order of dt to $\beta i_x dt$. Under the mean-field approximation, the expected state of individual x with k_x total social contacts after time dt can be written as $E\{s_x(t+dt)\} = s_x + (1-s_x)(\beta k_x \rho_t)dt$. Averaging this over all individuals, we get $\dot{\rho}_t = \rho_t(1-\rho_t)\beta\langle k\rangle$, where $\langle k\rangle$ is the average number of social contacts of individuals. Multiplying by the memory factor and integrating both sides, we get Equation (2).

Similar results can be derived for the voter-type alternative model discussed in the text. In this case, we have $E\{s_x(t+dt)\} = s_x + (1-s_x)\alpha dt(I_x/k_x)e^{-Bt}$, and if we sum this for all nodes, we get $\dot{\rho}_t = \rho_t(1-\rho_t)\alpha\sum_{xy}A_{xy}/kx$, where A_{xy} is the adjacency matrix of the underlying social network (it equals 1 if x is connected to y, and equals zero otherwise). Multiplying by the memory factor, the product of the two coefficients (α which pertains to social influence and $\sum_{xy}A_{xy}/kx$ which pertains to social connectivity) can be absorbed into a single new parameter, and after integration we arrive at a result identical to the one for the epidemic model.

Data Description

Film Data We extracted the publicly-available data set from www.boxofficemojo.com. We limited the analysis to the 5000 top US-grossing films. We excluded films before 1980 o ensure data reliability. We also excluded imax films which are not conventional feature films, and their lifetime is unusually long. For example, Space Station 3-D imax was released in 2002, and is still being screened. We limit the analysis only to 'traditional' films. We applied a cutoff of 70 weeks and omitted films with lifetimes longer than 70 weeks. Also, we excluded films with lifetimes shorter than 5 weeks. We also excluded the films released in 2016 to ensure that no film in the data set is still in theaters. The cleaned data set with sales time series accompanies this paper.

Wikipedia Data We extracted month-of-release Wikipedia page hits from the publicly-available data set provided by the Wikimedia foundation. Although the page hits data set goes back to 2007, we excluded films released prior to 2010 to make sure we are considering the period in which Wikipedia has become publicly established and popular enough to be a reliable source of data.

Author contributions statement

N.M and B.F. formulated the problem. A.T., N.M., and B.F. performed the research. A.T., N.M., B.F., and M.R. discussed the results and contributed to the text.

1 Competing Financial Interests Statement

The authors declare no competing financial interests.

References

1. T. Skocpol, *Social revolutions in the modern world*. Cambridge University Press, 1994.
2. T. Skocpol, *States and social revolutions: A comparative analysis of France, Russia and China*. Cambridge University Press, 1979.
3. H. Arendt, *Eichmann in jerusalem*. Penguin, 1963.
4. F. Neumann, H. Marcuse, and O. Kirchheimer, *Secret reports on Nazi Germany: The Frankfurt School contribution to the war effort*. Princeton University Press, 2013.
5. J. E. Stiglitz, *Freefall: America, free markets, and the sinking of the world economy*. WW Norton & Company, 2010.
6. R. G. Rajan, *Fault lines: How hidden fractures still threaten the world economy*. Princeton University Press, 2011.
7. S. Bikhchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of political Economy*, pp. 992–1026, 1992.
8. F. Davis, *Fashion, culture, and identity*. University of Chicago Press, 1994.
9. AV. Banerjee, "A simple model of herd behavior," *The Quarterly Journal of Economics*, vol. 107, no. 3, p. 797 1992.
10. D.S. Scharfstein and J.C. Stein, " Herd Behavior and Investment ," *The American Economic Review* , vol. 80, no. 3, p. 465 1990.
11. P.J. Peter, J.C. Olson, and K.G. Grunert, *Consumer behavior and marketing strategy*. McGraw-Hill London, 1999.
12. M.R. Solomon, *Consumer behavior: Buying, having, and being*. Prentice Hall Engelwood Cliffs, NJ, 2014.
13. D. J. Watts, *Everything is obvious:* Once you know the answer*. Crown Business, 2011.
14. J. Norton and F.M. Bass, "A diffusion theory model of adoption and substitution for successive generations of high-technology products," *Management Science*, vol. 33, no. 9, pp. 1069–1086, 1987.
15. R. Neelamegham and P. Chintagunta, "A bayesian model to forecast new product performance in domestic and international markets," *Marketing Science*, vol. 18, no. 2, pp. 115–136, 1999.
16. A. Ainslie, X. Drèze, and F. Zufryden, "Modeling movie life cycles and market share," *Marketing Science*, vol. 24, no. 3, pp. 508–517, 2005.
17. A. Elberse and J. Eliashberg, "Demand and supply dynamics for sequentially released products in international markets: The case of motion pictures," *Marketing Science*, vol. 22, no. 3, pp. 329–354, 2003.
18. J. Eliashberg, A. Elberse, and M. A. Leenders, "The motion picture industry: Critical issues in practice, current research, and new research directions," *Marketing science*, vol. 25, no. 6, pp. 638–661, 2006.
19. R. E. Krider and C. B. Weinberg, "Competitive dynamics and the introduction of new products: The motion picture timing game," *Journal of Marketing Research*, pp. 1–15, 1998.

20. L. Einav, "Seasonality in the us motion picture industry," *The Rand journal of economics*, vol. 38, no. 1, pp. 127–145, 2007.

21. S. A. Ravid, "Information, blockbusters, and stars: A study of the film industry*," *The Journal of Business*, vol. 72, no. 4, pp. 463–492, 1999.

22. M. Joshi, D. Das, K. Gimpel, and N. A. Smith, "Movie reviews and revenues: An experiment in text regression," in *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 293–296, Association for Computational Linguistics, 2010.

23. M. D. K. N. Smith, "Movie reviews and revenues: An experiment in text regression," *Boston Globe*, vol. 461, no. 154, p. 116.

24. S. Asur and B. A. Huberman, "Predicting the future with social media," in *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, vol. 1, pp. 492–499, IEEE, 2010.

25. K. R. Apala, M. Jose, S. Motnam, C.-C. Chan, K. J. Liszka, and F. de Gregorio, "Prediction of movies box office performance using social media," in *Advances in Social Networks Analysis and Mining (ASONAM), 2013 IEEE/ACM International Conference on*, pp. 1209–1214, IEEE, 2013.

26. W. Zhang and S. Skiena, "Improving movie gross prediction through news analysis," in *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 01*, pp. 301–304, IEEE Computer Society, 2009.

27. M. Mestyán, T. Yasseri, and J. Kertész, "Early prediction of movie box office success based on wikipedia activity big data," *PloS one*, vol. 8, no. 8, p. e71226, 2013.

28. S. Goel, J. M. Hofman, S. Lahaie, D. M. Pennock, and D. J. Watts, "Predicting consumer behavior with web search," *Proceedings of the National academy of sciences*, vol. 107, no. 41, pp. 17486–17490, 2010.

29. R. Pastor-Satorras, C. Castellano, P. Van Mieghem, and A. Vespignani, "Epidemic processes in complex networks," *Reviews of modern physics*, vol. 87, no. 3, p. 925, 2015.

30. S. Sudman and N. M. Bradburn, "Effects of time and memory factors on response in surveys," *Journal of the American Statistical Association*, vol. 68, no. 344, pp. 805–815, 1973.

31. Z. Lu, S. Williamson, and L. Kaufman, "Behavioral lifetime of human auditory sensory memory predicted by physiological measures," *Science* , vol. 258, pp. 1668–1668, 1992.

32. S. Sinha and S. Raghavendra, "Hollywood blockbusters and long-tailed distributions," *The European Physical Journal B-Condensed Matter and Complex Systems*, vol. 42, no. 2, pp. 293–296, 2004.

33. S. Sinha and R. K. Pan, "Blockbusters, bombs and sleepers: The income distribution of movies," in *Econophysics of Wealth Distributions*, pp. 43–47, Springer, 2005.

34. C. Castellano, S. Fortunato, and V. Loreto, "Statistical physics of social dynamics," *Reviews of modern physics*, vol. 81, no. 2, p. 591 2009.